

THE ELEMENTS OF MORAL PHILOSOPHY

Copyright © 1993, 1986 by McGraw-Hill, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

2 3 4 5 6 7 8 9 0 DOH DOH 9 0 9 8 7 6 5 4 3 2

ISBN 0-07-051098-9

This book was set in Baskerville by Arcata Graphics/Kingsport.
The editors were Cynthia Ward, Judith R. Cornwell, and Tom Holton;
the production supervisor was Friederich W. Schulte.
R. R. Donnelley & Sons Company was printer and binder.

Cover art: Alexander Rodchenko.

Non-Objective Painting: Black on Black. 1918

Oil on canvas, 32 1/4 x 31 1/4".

Collection, The Museum of Modern Art, New York, Gift of the artist,
through Jay Leyda. Photograph © The Museum of Modern Art, New York.

Library of Congress Cataloging-in-Publication Data

Rachels, James.

The elements of moral philosophy/James Rachels. — 2nd ed.

p. cm. — (The Heritage series in philosophy)

Includes bibliographical references and index.

ISBN 0-07-051098-9

1. Ethics. I. Title. II. Series.

BJ1012.R29

1993

92-12498

170—dc20

CHAPTER 9

Are There Absolute Moral Rules?

In spite of its horrifying title Kant's *Groundwork of the Metaphysics of Morals* is one of the small books which are truly great: it has exercised on human thought an influence almost ludicrously disproportionate to its size.

H. J. PATON, *THE MORAL LAW* (1948)

9.1. Kant and The Categorical Imperative

Imagine that someone is fleeing from a murderer and tells you he is going home to hide. Then the murderer comes along and asks where the first man went. You believe that if you tell the truth, the murderer will find his victim and kill him. What should you do—should you tell the truth or lie?

We might call this The Case of the Inquiring Murderer. In this case, most of us would think it is obvious what we should do: we should lie. Of course, we don't think we should go about lying as a general rule, but in these specific circumstances it seems the right thing to do. After all, we might say, which is more important, telling the truth or saving someone's life? Surely in a case such as *this* lying is justified.

There is one important philosopher, however, who thought we should *never* lie, even in a case such as this. Immanuel Kant (1724–1804) was one of the seminal figures in modern philosophy. Almost alone among the great thinkers, Kant believed that morality is a matter of following *absolute rules*—rules that admit no exceptions, that must be followed come what may. He believed, for example, that lying is never right, no matter what the circumstances. It is hard to see how

such a radical view could be defended, unless, perhaps, one held that such rules are God's unconditional commands. But Kant did not appeal to theological considerations; he relied only on rational arguments, holding that *reason* requires that we never lie. Let us see how he reached this remarkable conclusion. First we will look briefly at his general theory of ethics.

Kant observed that the word "ought" is often used non-morally. For example:

1. If you want to become a better chess player, you ought to study the games of Bobby Fischer.
2. If you want to go to law school, you ought to sign up to take the entrance examination.

Much of our conduct is governed by such "oughts." The pattern is: we have a certain wish (to become a better chess player, to go to law school); we recognize that a certain course of action would help us get what we want (studying Fischer's games, signing up for the entrance examination); and so we conclude that we should follow the indicated plan.

Kant called these "hypothetical imperatives" because they tell us what to do *provided that* we have the relevant desires. A person who did not want to improve his or her chess would have no reason to study Fischer's games; someone who did not want to go to law school would have no reason to take the entrance examination. Because the binding force of the "ought" depends on our having the relevant desire, we can *escape* its force simply by renouncing the desire. Thus by saying "I no longer want to go to law school," one can get out of the obligation to take the exam.

Moral obligations, by contrast, do not depend on our having particular desires. The form of a moral obligation is not "If you want so-and-so, then you ought to do such-and-such." Instead, moral requirements are *categorical*: they have the form, "You ought to do such-and-such, *period*." The moral rule is not, for example, that you ought to help people *if* you care for them or *if* you have some other purpose that helping them might serve. Instead, the rule is that you should be helpful to people *regardless of* your particular wants and desires. That is why, unlike hypothetical "oughts," moral requirements

cannot be escaped simply by saying "But I don't care about that."

Hypothetical "oughts" are easy to understand. They merely require us to adopt the means that are necessary to attain the ends we choose to seek. Categorical "oughts," on the other hand, are rather mysterious. How can we be obligated to behave in a certain way regardless of the ends we wish to achieve? Much of Kant's moral philosophy is an attempt to explain what categorical "oughts" are and how they are possible.

Kant holds that, just as hypothetical "oughts" are possible because we have desires, categorical "oughts" are possible because we have reason. Categorical "oughts" are binding on rational agents *simply because they are rational*. How can this be so? It is, Kant says, because categorical oughts are derived from a principle that every rational person must accept. He calls this principle *The Categorical Imperative*. In his *Groundwork of the Metaphysics of Morals* (1785), he expresses The Categorical Imperative like this:

Act only according to that maxim by which you can at the same time will that it should become a universal law.

This principle summarizes a procedure for deciding whether an act is morally permissible. When you are contemplating doing a particular action, you are to ask what rule you would be following if you were to do that action. (This will be the "maxim" of the act.) Then you are to ask whether you would be willing for that rule to be followed by everyone all the time. (That would make it a "universal law" in the relevant sense.) If so, the rule may be followed, and the act is permissible. However, if you would *not* be willing for everyone to follow the rule, then you may not follow it, and the act is morally impermissible.

Ex. Kant gives several examples to explain how this works. Suppose, he says, a man needs to borrow money, and he knows that no one will lend it to him unless he promises to repay. But he also knows that he will be unable to repay. He therefore faces this question: Should he promise to repay the debt, knowing that he cannot do so, in order to persuade someone to make the loan? If he were to do that, the "maxim

of the act" (the rule he would be following) would be: *Whenever you need a loan, promise to repay it, even though you know you cannot do so.* Now, could this rule become a universal law? Obviously not, because it would be self-defeating. Once this became a universal practice, no one would any longer believe such promises, and so no one would make loans because of them. As Kant himself puts it, "no one would believe what was promised to him but would only laugh at any such assertion as vain pretense."

EX. Another of Kant's examples has to do with giving charity. Suppose, he says, someone refuses to help others in need, saying to himself "What concern of mine is it? Let each one be happy as heaven wills, or as he can make himself; I will not take anything from him or even envy him; but to his welfare or to his assistance in time of need I have no desire to contribute." This, again, is a rule that one cannot will to be a universal law. For at some time in the future this man might *himself* be in need of assistance from others, and he would not want others to be so indifferent to him.

9.2. Absolute Rules and the Duty Not to Lie

Being a moral agent, then, means guiding one's conduct by "universal laws"—moral rules that hold, without exception, in all circumstances. Kant thought that the rule against lying was one such rule. Of course, this was not the *only* absolute rule Kant defended—he thought there are many others; morality is full of them. But it will be useful to focus on the rule against lying as a convenient example. Kant devoted considerable space to discussing this rule, and it is clear that he felt especially strongly about it—he said that lying in any circumstances is "the obliteration of one's dignity as a human being."

Kant offered two main arguments for this view. Let us examine them one at a time.

1. His primary reason for thinking that lying is always wrong was that the prohibition of lying follows straightaway from The Categorical Imperative. We could not will that it be a universal law that we should lie, because it would be self-defeating; people would quickly learn that they could not rely on what other people said, and so the lies would not be be-

lieved. Surely there is something to this: in order for a lie to be successful, people must believe that others are telling the truth; so the success of a lie depends on there *not* being a "universal law" permitting it.

There is, however, an important problem with this argument, which will become clear if we spell out Kant's line of thought more fully. Let us return to The Case of the Inquiring Murderer. Should you tell him the truth? Kant would have you reason as follows:

- ✱ (1) You should do only those actions that conform to rules that you could will to be adopted universally.
- (2) If you were to lie, you would be following the rule "It is permissible to lie."
- (3) This rule could not be adopted universally, because it would be self-defeating: people would stop believing one another, and then it would do no good to lie.
- (4) Therefore, you should not lie.

The problem with this way of reasoning was nicely summarized by the British philosopher Elizabeth Anscombe when she wrote about Kant in the academic journal *Philosophy* in 1958:

His own rigoristic convictions on the subject of lying were so intense that it never occurred to him that a lie could be relevantly described as anything but just a lie (e.g., as "a lie in such-and-such circumstances"). His rule about universalizable maxims is useless without stipulations as to what shall count as a relevant description of an action with a view to constructing a maxim about it.

The difficulty arises in step (2) of the argument. Exactly what rule would you be following if you lied? The crucial point is that there are many ways to formulate the rule; some of them might not be "universalizable" in Kant's sense, but some would be. Suppose we said you were following *this* rule (R): "It is permissible to lie when doing so would save someone's life." We *could* will that (R) be made a "universal law," and it would not be self-defeating.

It might be replied that the universal adoption of (R) *would* be self-defeating because potential murderers would cease to believe us. But they would believe us if they thought we did not know what they were up to; and if they thought we *did* know what they were up to, they would not bother to ask us in the first place. This is no different from the situation that exists now, in the real world: murderers know that people will not willingly aid them. Thus the adoption of (R) would help save lives, at little cost, and it would not undermine general confidence in what people say in ordinary circumstances.

2. The Case of the Inquiring Murderer is not simply an example I made up; it is Kant's own example. In an essay with the charmingly old-fashioned title "On a Supposed Right to Lie from Altruistic Motives," Kant discusses this case and gives a second argument for his view about it. He writes:

After you have honestly answered the murderer's question as to whether his intended victim is at home, it may be that he has slipped out so that he does not come in the way of the murderer, and thus that the murder may not be committed. But if you had lied and said he was not at home when he had really gone out without your knowing it, and if the murderer had then met him as he went away and murdered him, you might justly be accused as the cause of his death. For if you had told the truth as far as you knew it, perhaps the murderer might have been apprehended by the neighbors while he searched the house and thus the deed might have been prevented. Therefore, whoever tells a lie, however well intentioned he might be, must answer for the consequences, however unforeseeable they were, and pay the penalty for them. . . .

To be truthful (honest) in all deliberations, therefore, is a sacred and absolutely commanding decree of reason, limited by no expediency.

This argument may be stated in a more general form: We are tempted to make exceptions to the rule against lying because in some cases we think the consequences of truthfulness would be bad and the consequences of lying good. However, we can never be certain about what the consequences of our actions will be; we cannot *know* that good results will follow.

The results of lying *might* be unexpectedly bad. Therefore, the best policy is always to avoid the known evil—lying—and let the consequences come as they will. Even if the consequences are bad, they will not be our fault, for we will have done our duty.

The problems with this argument are obvious enough—so obvious, in fact, that it is surprising a philosopher of Kant's stature was not more sensitive to them. In the first place, the argument depends on an unreasonably pessimistic view of what we can know. Sometimes we can be quite confident of what the consequences of our actions will be, and justifiably so; in which case we need not hesitate because of uncertainty. Moreover—and this is a more interesting matter, from a philosophical point of view—Kant seems to assume that although we would be morally responsible for any bad consequences of lying, we would *not* be similarly responsible for any bad consequences of telling the truth. Suppose, as a result of our telling the truth, the murderer found his victim and killed him. Kant seems to assume that we would be blameless. But can we escape responsibility so easily? After all, we aided the murderer. This argument, then, like the first one, is not very convincing.

9.3. Conflicts Between Rules

The idea that moral rules are absolute, allowing no exceptions, is implausible in light of such cases as The Case of the Inquiring Murderer, and Kant's arguments for it are unsatisfactory. But are there any convincing arguments against the idea, apart from its being implausible?

The principal argument against absolute moral rules has to do with the possibility of conflict cases. Suppose it is held to be absolutely wrong to do A in any circumstances and also wrong to do B in any circumstances. Then what about the case in which a person is faced with the choice between doing A and doing B—when he must do something and there are no other alternatives available? This kind of conflict case seems to show that it is *logically* untenable to hold that moral rules are absolute.

Is there any way that this objection can be met? One way would be for the absolutist to deny that such cases ever actu-

ally occur. The British philosopher P. T. Geach takes just this view. Like Kant, Geach argues that moral rules are absolute; but his reasons are very different from Kant's. Geach holds that moral rules must be understood as absolute divine commands, and so he says simply that God will not allow conflict situations to arise. We can describe fictitious cases in which there is no way to avoid violating one of the absolute rules, but, he says, God will not permit such circumstances to exist in the real world. In his book *God and the Soul* (1969) Geach writes:

"But suppose circumstances are such that observance of one Divine law, say the law against lying, involves breach of some other absolute Divine prohibition?"—If God is rational, he does not command the impossible; if God governs all events by his providence, he can see to it that *circumstances in which a man is inculpably faced by a choice between forbidden acts do not occur*. Of course such circumstances (with the clause "and there is no way out" written into their description) are consistently describable; but God's providence could ensure that they do not in fact arise. Contrary to what nonbelievers often say, belief in the existence of God does make a difference to what one expects to happen.

Do such circumstances ever actually arise? The Case of the Inquiring Murderer is, of course, a fictitious example; but it is not difficult to find real-life examples that make the same point. During the Second World War, Dutch fishermen regularly smuggled Jewish refugees to England in their boats, and the following sort of thing sometimes happened. A Dutch boat, with refugees in the hold, would be stopped by a Nazi patrol boat. The Nazi captain would call out and ask the Dutch captain where he was bound, who was on board, and so forth. The fishermen would lie and be allowed to pass. Now it is clear that the fishermen had only two alternatives, to lie or to allow their passengers (and themselves) to be taken and shot. No third alternative was available; they could not, for example, remain silent and outrun the Nazis.

Now suppose the two rules "It is wrong to lie" and "It is wrong to permit the murder of innocent people" are both taken to be absolute. The Dutch fishermen would have to do

one of these things; therefore a moral view that absolutely prohibits both is incoherent. Of course this difficulty could be avoided if one held that only *one* of these rules is absolute; that would apparently be Kant's way out. But this dodge cannot work in every such case; so long as there are at least two "absolute rules," whatever they might be, the possibility will always exist that they might come into conflict. And that makes the view of those rules as absolute impossible to maintain.

9.4. Another Look at Kant's Basic Idea

Few philosophers would dispute Paton's statement that Kant's *Groundwork* "has exercised on human thought an influence almost ludicrously disproportionate to its size." Yet at the same time, few would defend The Categorical Imperative as Kant formulated it—as we have seen, it is beset by serious, perhaps insurmountable, problems. What, then, accounts for Kant's influence? Is there some basic idea underlying The Categorical Imperative that we might accept, even if we do not accept Kant's particular way of expressing it? I believe that there is, and that the power of this idea accounts, at least in part, for Kant's vast influence.

Remember that Kant thinks The Categorical Imperative is binding on rational agents simply because they are rational—in other words, a person who did not accept this principle would be guilty not merely of being immoral but of being *irrational*. This is a fascinating idea—that there are rational as well as moral constraints on what a good person may believe and do. But what exactly does this mean? In what sense would it be irrational to reject The Categorical Imperative?

The basic idea seems to be this: A moral judgment must be backed by good reasons—if it is true that you ought (or ought not) to do such-and-such, then there must be a *reason why* you should (or should not) do it. For example, you may think that you ought not to set forest fires because property would be destroyed and people would be killed. But if you accept those as reasons in *one* case, you must also accept them as reasons in *other* cases. It is no good saying that you accept those reasons some of the time, but not all the time; or that other people must respect them, but not you. Moral reasons,

if they are valid at all, are binding on all people at all times. This is a requirement of consistency; and Kant was right to think that no rational person could deny it.

This is the Kantian idea—or, I should say, one of the Kantian ideas—that has been so influential. It has a number of important implications. It implies that a person cannot regard himself as special, from a moral point of view: he cannot consistently think that *he* is permitted to act in ways that are forbidden to others, or that *his* interests are more important than other people's interests. As one commentator remarked, I cannot say that it is all right for me to drink your beer and then complain when you drink mine. Moreover, it implies that there are *rational constraints* on what we may do: we may want to do something—say, drink someone else's beer—but recognize that we cannot *consistently* do it, because we cannot at the same time accept its implications. If Kant was not the first to recognize this, he was the first to make it the cornerstone of a fully worked-out system of morals. That was his great contribution.

But Kant went one step further and concluded that consistency requires rules that have no exceptions. It is not hard to see how his basic idea pushed him in that direction; but the extra step was not necessary, and it has caused trouble for his theory ever since. Rules, even within a Kantian framework, *need not* be regarded as absolute. All that is required by Kant's basic idea is that when we violate a rule, we do so for a reason that we would be willing for anyone to accept, were they in our position. In *The Case of the Inquiring Murderer*, this means that we may violate the rule against lying only if we would be willing for anyone to do so were he faced with the same situation. And *that* proposition causes little trouble.